PRTN 174

Packet Radio Network
Station Labeling Process

Julie Sussman, BBN

1. Introduction

This memo discusses the process in the station that is responsible for network labeling, concentrating on the initial version now being implemented. It does not attempt to specify the final station labeling procedures, and does not even touch on many important issues. Future revisions of this memo will be more complete.

2. The Network

This section briefly describes the network being implemented, concentrating on those aspects of particular interest for labeling purposes.

2.1 PRs

The Packet Radio Network (PRN) consists of Packet Radio Units (PRs) either standing alone (repeaters) or attached to a terminal or station (terminal PRs and station PRs). PRs receive and transmit packets of information directly from/to their attached terminal or station, if any, and via radio to other PRs. The function of a terminal or station PR is to transmit packets between the network and its attached device; the function of a repeater PR is to relay packets between other PRs. Every PR has a unique built-in ID, coded so that a PR's function (repeater, terminal PR, station PR) can be determined from its ID.

2.2 Labels

The network is organized for packet routing into hierarchy levels. Every PR is assigned to a level (the number of PRs a packet must pass through in traveling between the PR and the station) and given a label unique within that level. When a PR is initialized, it is unlabeled. The station can also unlabel a (labeled) PR by sending it an appropriate packet. Unlabeled PRs cannot participate in packet transport. A PR becomes labeled upon receiving a label packet from the station addressed to it (containing its ID), assigning it a level, label, and route to the station. The route consists of a label for each level between the PR and the station PR inclusive. (If the recommendations of PRTN 162 are adopted, the route may be incomplete -- i.e., it may not contain labels for every level.) The station also tells the PR how to locate the route fields it will need to look at -- in particular, the field for its level; knowledge of route format is not built into the PR.

2.3 Routing

Every packet contains a route in its header defining the label of the PR at each level that it should pass through. The packet is transmitted from the PR that has it to the PR at the next level whose label is next in the packet route, until it reaches its destination. If a PR does not hear a hop-by-hop acknowledgement (HBH ACK) for a packet it transmits, it may retransmit, or eventually alternate route the packet by re-addressing it to any PR at the next level. If a packet departs from its original route, its route is modified, so that upon receipt a packet's header always contains the route it has actually traversed.

2.4 ROPs

The purpose of a ROP (Repeater On Packet) is to inform the station of the existence and labeling of a PR and what PRs can hear its transmissions. Every PR generates ROPs at regular intervals (though it will delay the ROP if it is busy), incrementing the sequence number for each ROP. Unlabeled PRs generate ROPs much more frequently than labeled PRs. The text of the ROP tells whether the PR is labeled and, if so, what its level, label and route are. The initial transmission of a ROP differs from that of other packets in that the originator only transmits it once, not expecting an acknowledgement, and any labeled PR that receives it will accept it regardless of its level or label. This PR then inserts its own ID into the ROP and forwards it along its own route to the station. From this point on, the ROP follows normal packet transport protocol. NB: The station PR transmits its ROPs directly to the attached station as well as via radio.

2.5 Reliability

The source of a packet may request that the destination acknowledge receipt of it. In this case, the destination sends back the packet header as an ETE ACK (end-to-end acknowledgement). The source may retransmit until it receives the ACK.

2.6 Characteristics

Two important characteristics of packet transmission in the PRN are:

1) Packets from a single source may arrive at their destination out of order -- i.e., not in the order in which they were sent.

2) Due to retransmission and alternate routing, duplicate copies of a packet are frequently generated. Although some copies may be filtered out en route, more than one may arrive at the destination.

3. The Labeling Process

3.1 Design goals

The station is responsible for labeling the network. It should provide working labels (i.e., level/label/route) to all PRs so as to create good traffic flow. More specifically, it should:

- Label unlabeled PRs.

- Relabel labeled PRs whose labels are not longer legitimate. A label may become bad either because the station is reinitialized, in which case all PRs must be relabeled, or because some link (two-way communication path between two PRs) on a route no longer works. Some reasons for link failure are: changed radio conditions stop communication between still-operating PRs; a mobile terminal moves out of range of PRs it used to be able to communicate with; a PR stops operating.

- Relabel labeled PRs so as to improve traffic flow.

Assigned routes should always work in both directions, since the hop acknowledgement retransmission protocol transmits the packet in both directions over each hop.

3.2 Design considerations

What information is available to the station?

What can the station deduce about network conditions from this information?

What actions will the station take in various situations?

How are these actions accomplished?

3.3 Labeling process - 1st version

3.3.1 Overview

The first version of the labeling process (being implemented now) is quite simple and limited in capability. It does not perform all the functions specified above in 3.1.

It does:

give unlabeled PRs labels which are likely to work;

relabel PRs operating with old labels when the station is re-initialized;

relabel PRs if some PR on their route becomes unlabeled.

It does not:

check that assigned routes work (and continue to work) -- in particular, it does not track mobile terminals or notice failure of a repeater;

attempt to improve traffic flow.

The program uses the ROPs it receives from PRs as its only source of information; these tell it about the existence, labeling, and connectivity of PRs. In fact, it is entirely ROP-driven; it only labels (or unlabels) a PR when stimulated to do so by receipt of a ROP from it. The only actions taken by the program are labeling and unlabeling PRs.

3.3.2 Deducing and changing labels

When the station receives a ROP from a PR, telling it what label (if any) the PR is operating with, it decides whether to change the PR's labeling as follows:

-If the PR is using a label the station didn't give it (ROP label is not recorded in station tables), it will be unlabeled. (This takes care of the case where the station is re-initialized in an already-labeled network, as well as any other inconsistencies which may arise.)

-If the PR is unlabeled and a label can be chosen for it (see next section), it will be labeled. (The station may not actually label a PR immediately upon learning that it is unlabeled -- it may wait until a later ROP, so as to have more information on which to base label selection.)

-If the PR is marked in the station tables as needing relabeling and a label can be chosen for it (see next section), it is relabeled. Currently a PR is only marked for relabeling when the labeling of some PR on its assigned route changes.

To unlabel a PR, the appropriate command (not yet specified by Collins, but probably a control packet -- i.e., a packet containing code for the PR to execute) is sent to it; to label a PR, a label packet is sent to it routed along the route to be assigned. (The actual packet sending, including retransmission and recognition of acknowledgement, is handled by the station connection process, which interfaces all station application processes to the network.) No further commands are sent to the PR until the outcome of this one is known for sure, regardless of receipt of further ROPs.

Digitized by the Internet Archive
in 2025 with funding from
Amateur Radio Digital Communications, Grant 151

https://archive.org/details/temporary-note-174

If an ACK for the label/unlabel packet is received, then the PR received the packet; its new labeling is recorded. The station must be careful, however, in interpreting ROPs from the PR for a while. Because the network may reorder packets, it is possible for ROPs generated by the PR before the ACK (i.e., before it received the new label or unlabel) to arrive after it. Such ROPs contain obsolete label information. Thus label information from ROPs received after the ACK is ignored for a period equal to the maximum lifetime of a packet in the network; i.e., these ROPs are not used as a basis for deciding to change the PR's labeling.

If no ACK is received after several transmissions, the station gives up trying to send the packet. It does not know at this point whether or not the label (or unlabel) got to the PR. It cannot immediately conclude from ROPs coming in with the PR's old status that the packet didn't arrive, since the ROPs could have been generated before its receipt. Thus, as in the ACK case above, the station waits a while before believing the status information in the PR's ROPs to be definitive. In this case, it waits for a period equal to the maximum round trip time of a packet (the maximum lifetime of a packet and its ACK) after it last transmitted the command. Any ROPs received after this time must have been generated after the last possible time the PR could have received the command; hence they contain its true status.

3.3.3 Choosing labels

The labeling program needs information on the radio connectivity of the network in order to assign valid labels and routes to PRs. Since the radio link between two PRs may be operational in only one direction, the program keeps track, for each PR, of whether it can hear each other PR. One of three possible values is recorded for each ordered pair of PRs: YES (a transmission has been heard over this hop), NO (this hop has been tried and found wanting), and MAYBE (neither YES nor NO). The reason for three values instead of two (YES/NO) will be seen below.

When a new PR is entered into the station tables, all of its connectivity is initialized to MAYBE.

The connectivity for each hop traversed by every ROP received is set to YES. The first hop is identified from the two IDs in the packet, those of the PR that originated the ROp and the PR that heard and relayed it; the rest are deduced from the labels in the packet route.

[In addition, since all packets entering the station pass through the connection process, the connection process could set connectivity to YES for hops traversed by other packets. It might record connectivity from the route of every packet received, or some subset of arriving packets -- for example:

ACKs; packets destined for station processes (i.e., not being forwarded); every nth packet; etc. The decision on what packets to set connectivity from will depend on the effect of this processing on station performance.]

Whenever a connectivity entry is modified, a flag corresponding to that entry is set. At regular intervals, all entries that haven't been modified (since the last time this was done) are re-initialized to MAYBE, thus erasing any old YES/NO information, and the "modified" flags are reset.

To choose a label for a PR (call it P), the program looks for a PR to route its traffic through. It searches its tables for a PR which
-can hear P (connectivity from P to it is YES)
-can or may be able to be heard by P (connectivity from it to P is YES or MAYBE)
-is not at the highest possible level (since P will be placed a level higher than it); also there must be an available label value at the next higher level
-is not a terminal PR (they don't forward packets for other PRs)
-is believed to be stably labeled (is labeled and is not in the process of having its label changed or marked as needing relabeling)

(The reason for not insisting on YES connectivity from the PR to P is that an unlabeled PR doesn't forward packets it receives, so no connectivity to it is known; however, connectivity from it is known from its ROPs.)

If there is more than one such PR, preference is given to one which
-P can definitely hear (YES, not MAYBE - so route is more likely to work)
-has a lower level (so the number of hops between P and the station will be minimized)
-has fewer terminals already routing through it (a rudimentary attempt to spread traffic out)

(NB: The order of listing criteria is not meant to imply an order of applying them.)

If such a PR is found, P is assigned the next higher level, a label not in use at that level, and a route consisting of that PR's route with the hop between it and P appended. (If the recommendations of PRTN 162 are adopted, only the single hop between the PR and P will actually be assigned.)

If the program gives up trying to deliver a label (no ACK is received), it sets the connectivity for the last outbound hop (the hop to the PR being labeled) in the label it tried to assign to NO. It of course doesn't know what hop's failure prevented

the label's delivery; it might have been an earlier hop, or perhaps the label actually was delivered but the ACK couldn't get in from the PR to the station. Blaming the last outbound hop is as good a guess as any, and in some cases an excellent guess, since everything up to the last hop was believed to be working when the label was chosen, but that hop may have been in doubt. In any case, setting this hop to NO prevents choosing the same route again -- at least until connectivity is refreshed (old values reset to MAYBE).

3.3.4 Route format

The allocation of bits in a route among the various levels determines how many labels, hence how many PRs, can exist at each level. A default route format is built into the station, but the operator can modify the format during station initialization to be appropriate for the expected network configuration. This format remains fixed throughout the life of the station.

> [Collins has proposed that the station labeling process figure out for itself how to allocate the fields, probably based on observing ROPs, and change the allocation dynamically as needed. Methods for doing this have not been worked out, and we leave their consideration for the future.]

3.4 Coming attractions

The next version of the labeling process will include some attempt to check that assigned routes work. This will probably involve the use of probe packets (executed by the PR) which return to the station a packet telling what route the probe actually followed.